

Incorporating Mental Simulation for a More Effective Robotic Teammate

**William G. Kennedy, Magdalena D. Bugajska, William Adams,
Alan C. Schultz, J. Gregory Trafton**

Naval Research Laboratory, 4555 Overlook Ave. SW, Washington, DC 20375
[Bill.Kennedy, Magda.Bugajska, William.Adams, Alan.Schultz, Greg.Trafton]@NRL.Navy.mil

Abstract

How can we facilitate human-robot teamwork? The teamwork literature has identified the need to know the capabilities of teammates. How can we integrate the knowledge of another agent's capabilities for a justifiably intelligent teammate? This paper describes extensions to the cognitive architecture, ACT-R, and the use of artificial intelligence (AI) and cognitive science approaches to produce a more cognitively-plausible, autonomous robotic system that "mentally" simulates the decision-making of its teammate. The extensions to ACT-R added capabilities to interact with the real world through the robot's sensors and effectors and simulate the decision-making of its teammate. The AI applications provided visual sensor capabilities by methods clearly different than those used by humans. The integration of these approaches into intelligent team-based behavior is demonstrated on a mobile robot. Our "TeamBot" matches the descriptive work and theories on human teamwork. We illustrate our approach in a spatial, team-oriented task of a guard force responding appropriately to an alarm condition that requires the human and robot team to "man" two guard stations as soon as possible after the alarm.

Introduction

What is the best way to build robots that interact with humans as teammates? Our approach uses a novel mixture of AI, cognitive science, and robotics research to build robotic teammates that work well with human teammates. This paper describes a project that integrates advances in cognitive modeling, AI, and robotics to demonstrate improvement in the behavior of an autonomous robot acting appropriately as a part of a human-robot team providing warehouse security.

Teams and Teamwork

Teams and teamwork, including multi-agent collaboration, are large research areas within both artificial intelligence and cognitive science. Even a brief discussion of relevant papers within both fields would require a long survey article in itself. Here, we mention only the most seminal.

This work is not subject to U.S. copyright protection. Copyright © 2008 all other countries, Association for the Advancement of Artificial Intelligence.

Teamwork: the AI Perspective

General models of teamwork and collaboration within AI include: STEAM and TEAMCORE (Tambe 1997), SharedPlans (Grosz and Kraus 1996), COLLAGEN (Rich and Sidner 1996), and RETSINA (Sycara et al. 2003). For a broad overview of teamwork in multi-agent systems, see (Stone and Veloso 2000 and Sycara 1998), and in particular, (Sycara and Sukthankar 2006).

Key issues in multi-agent systems research include the organization and make up of teams, task allocation among team members, multi-agent planning (including recognizing and resolving conflicts among agents and within plans), managing limited resources, communications among agents (including contingencies for when there is no communication), adaptation and learning in the team, and agent tracking and monitoring.

Of particular importance to our interests are systems that allow agents to learn about and model their *own* teammates and then use that knowledge to improve collaboration. (Kaminka et al. 2003) presents a technique that allows one agent (a coach) to predict the future behavior of other agents (its own team and the opponent team) in order to coordinate activities by observing those agents and building a model of their behavior. Observations are translated into a time series of recognized atomic behaviors, and these into subsequences that characterize a team (although not necessarily a single agent). (Kaminka and Tambe 2000) investigated just how much monitoring of another agent is sufficient for an agent to be an effective teammate.

Our approach is to model the other agent in order to reduce the amount of monitoring that is required, and doing so in a cognitively plausible way by having the robotic team member perform a mental simulation, based on its computational cognitive model of the other agent.

Teamwork: the Psychology Perspective

There are many studies on what makes an effective team (Cannon-Bowers, Salas, and Converse 1993; Cooke, et al. 2000; Mathieu et al. 2000; Stout, Cannon-Bowers, Salas, and Milanovich 1999). This research addresses the same key issues as the AI research. Their methodology also includes examining how high and low performing teams accomplish team-related tasks. The research suggests that

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2008		2. REPORT TYPE		3. DATES COVERED 00-00-2008 to 00-00-2008	
4. TITLE AND SUBTITLE Incorporating Mental Simulation for a More Effective Robotic Teammate				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory, 4555 Overlook Ave. SW, Washington, DC, 20375				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 6	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

the knowledge employed by a good team member has three components (Cannon-Bowers, Salas, and Converse 1993):

- (1) Knowledge of own capabilities [meta-knowledge],
- (2) Knowledge of the task, and
- (3) Knowledge about the capabilities of their teammates.

Most researchers have suggested that these three components are deeply inter-related; i.e., without any one of these, a person is not a good team member.

We believe the first two have been addressed if the agent is competent in the task domain as an individual. Here, we address the third component, the knowledge of a teammate's capabilities, focusing on the teammate's cognitive processes. The sharing of this understanding among teammates is frequently called a shared mental model and has been suggested to be key to understanding team performance (Cannon-Bowers, Salas, and Converse 1993).

Teamwork: Mental Simulation

It appears the capability to model or simulate others is part of human cognition. There is evidence that infants become capable of predicting the actions of others in the second half of their first year as they develop the capability to perform those actions themselves (Falak-Ytter, Gredeback, and Von Hofsten 2006). This suggests that the ability to simulate others is inherent or developed very early. Research on teamwork has reported that human team performance distributions differed significantly but the distributions did not converge over time (Mathieu et al. 2000). We interpret that to be further evidence that the capability to build mental models of teammates was not developed or learned during the tests, but must be already available. Therefore, it may be appropriate for the capability to be part of the cognitive modeling system. We based the work described in this paper on the premise that humans model others and use themselves, i.e., their own capabilities, as an initial model of their teammate.

Simulating the cognitive processes of another agent requires maintaining multiple states of reality. Many current cognitive architectures do not have such a mechanism. The problem spaces in Soar (Rosenbloom, Laird and Newell 1993) and alternate worlds in Polyscheme (Cassimatis 2005) are the best-known exceptions. Soar's problem spaces facilitate subgoaling and have been used to anticipate opponent's behavior in the game of Quake (Laird 2001). Polyscheme's worlds are a general construct and allow for instantiation and manipulation of hypothetical, counterfactual, and even stochastic simulations. The alternate worlds in Polyscheme have been used to model spatial perspective-taking (Cassimatis et al. 2004) and theory of mind (Bello and Cassimatis 2006).

The concept of simulating the cognitive processes of another agent ties to previous work (Trafton et al. 2005) where it was suggested that an important consideration in designing an architecture for integrated intelligence, is how well the system works with a person. As a result, we were

guided by *the representation hypothesis* that suggests that when a system uses representations and processes similar to a person's, it will be able to collaborate with a person better than a computational system that does not. Furthermore, such a system will be more compatible with human expectations of reasonable behavior, and thus more accommodating to the human.

Our principal goal in this project is to show how the integration of mental simulation of a teammate within an embodied computational cognitive model can improve performance of the robotic teammate.

Laboratory Scenario

As a test bed for our research, we introduce a task for a human-robot team that demonstrates the usefulness of the ability to model the decision-making of a teammate. The scenario takes place in a specified area, such as a warehouse, where teammates can frequently see each other and move freely through the area. In this area, a robot and a human are a security team charged with patrolling the area and responding to alarms. While the human and robot are separately patrolling inside the area an alarm sounds. Their task is then to 'man' the two widely separated guard stations as soon as possible after the alarm.

To successfully complete this task, both team members must be inside different guard stations. If both go to the same guard station, there is a cost in the team's response time for one of them to go to the other station. We did not allow communication between them during the response. For the robot, the task requires it to know the whereabouts of its teammate while it is patrolling to be prepared for the alarm, know the spatial locations of the guard stations, be able to perform spatial reasoning to judge the distances to the guard stations (specifically which station is closer), and predict where the human will go based on spatial reasoning. Figure 1 is a diagram of the task.




Figure 1. Security team's alarm response decisions

Approach

This project extends recent work (Trafton et al. 2008) by developing ACT-R/E. ACT-R/E (embodied) is based on ACT-R (Anderson 2007; Anderson et al. 2004) and is an embodied system that thinks and reasons like people do.

The ACT family of theories has a long history of successfully matching psychological process data and has been broadly tested in psychological and computational terms.

ACT-R/E

ACT-R is a hybrid symbolic/sub-symbolic production-based system. It consists of a number of modules and their interfaces and buffers. Modules in ACT-R are intended to represent relatively specific cognitive faculties such as declarative (fact-based) and procedural (rule-based) memory, visual and auditory perception, vocalization, and time perception. Buffers in ACT-R make up the working memory of a cognitive model. Some modules fill their buffers in response to the changes in the environment and all modules fill their buffers in response to explicit procedural requests. Like many production systems, ACT-R continuously matches production conditions against the working memory (buffers), selects a single production to fire, and then executes specified buffer changes and module requests which eventually result in updates to relevant buffers.

This project embodied ACT-R on an iRobot B21r, a human-scale, zero-turn-radius robotic platform suited to use in indoor environments. It carries the sensors and provides onboard computing support for all sensing, navigation, and output. With ACT-R/E, we have extended the ACT-R architecture with rudimentary spatial reasoning (spatial module), localization and navigation faculties (“moval” module), and modified the visual, aural, and vocal modules to use actual robot sensors as shown in the architectural diagram in Figure 2. The modules will be discussed with the behaviors they provide.

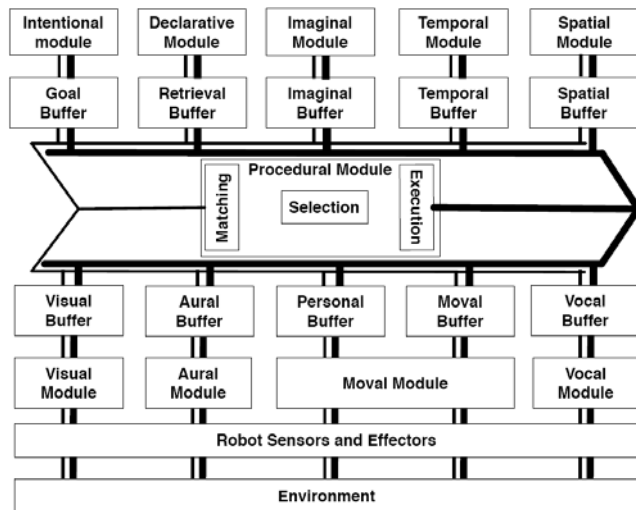


Figure 2. ACT-R/E architecture.

Modeling Task Behavior

The model of an agent’s behavior in the warehouse guard force task was achieved by modeling four distinct behaviors.

Patrolling the perimeter. Based on the localization information obtained from the personal buffer, the agent is capable of retrieving from declarative memory the next waypoint on the perimeter and issuing a request to the moval module to navigate to that point. The moval module works through the mobility system that provides localization, collision avoidance, path planning, and map learning (Schultz, Adams, and Yamauchi 1999).

Maintaining spatial situation awareness. In order to successfully perform this task, the agent needs to be aware of its own spatial location as well as that of its team member and the points of interest, which in this scenario are the security stations. The spatial module is capable of the necessary spatial reasoning. During the task, a cognitive map continuously maintains the robot’s own location, but updates to the location of the teammate are only made in response to explicit procedural requests that are based on visual observations. The visual module uses a person-tracking system (Fransen et al. 2007) that exploits an omni-directional camera mounted on top of the robot. Cognitive plausibility is addressed by limiting this sensor to a human’s nominal 140-degree field of view. The robotic agent’s visual focus of attention is indicated by rotating the animated face (Parke and Waters 1996; Simmons et al. 2003) displayed on the robot-mounted LCD monitor to face the object being attended to.

Listening for the alarm. While patrolling, the system continuously monitors the aural buffer for the alarm event and in response to a detection, it instantiates the alarm response behavior. Currently, the sound detection capabilities of the aural system (Fransen et al. 2007; Trafton et al. 2008), are limited to recognizing the phrase “alarm” and detecting certain relevant sounds within predetermined frequency bands such as human speech and alarm sirens. The detection of the alarm is verbally reported by issuing a request to the vocal module which implements speech using the Cepstral Swift speech generation system.

Responding to the alarm. In response to the alarm, the agent chooses a security station to “man,” retrieves that location from declarative memory, and issues a request to the moval module to navigate to the recalled location. If the agent detects the teammate in the same station (a “conflict”), it proceeds to the other location to complete the task. Currently, the conflict detection is based solely on sight, but could involve any form of perception or explicit communication.

In this work, we have modeled the following strategies for choosing the station:

- (1) Self-centered strategy. The agent determines the station closest to its own position and goes there.
- (2) Collaborative strategy. The agent predicts the teammate’s choice of destination and goes to the *other* station in an attempt to avoid conflicts.

The challenge for our work was the modeling of the second strategy, a response based on simulating a teammate’s cognitive processes.

Mental Simulation in ACT-R/E

To enable TeamBot to simulate its teammate's behavior, we used a capability of the ACT-R architecture which facilitates running additional cognitive models simultaneously. An ACT-R model basically consists of declarative and procedural memory and an initial goal. Similar to Soar and Polyscheme implementations, the ability of ACT-R to spawn a new model from within a running model allows us to represent and manipulate a mental model of another agent. To allow the base cognitive model to continue running while the simulation occurs, we run the two models synchronously at the production-level. The flexibility of fixing the declarative memory and productions of the simulated mental model to a subset of the original model's allows the system to consider hypothetical and counterfactual situations.

Modeling the Teammate's Decision Making

We start with the premise that humans use themselves as an approximate, initial model of their teammate. Therefore, we based the simulation of the teammate's decision making on the robot's own knowledge of the situation and its decision process. To predict the teammate's choice of destination in response to the alarm for the collaborative strategy, we modeled the human as following the self-centered strategy, i.e., that the human will go to the closest station. The result of the simulation, i.e., which station the human is expected to go to, is made available to the base model by inserting the result into the "imaginal" buffer. The availability of the results of the mental simulation facilitates the agent's completion of its own decision making. The effect is that the robotic agent yields to what it believes is the human's choice.

While this simple model of teamwork allows us to demonstrate the concept and the implementation of the simulation of the teammate, it should be noted that the simulation could have used the collaborative strategy recursively. In other words, we could have modeled the interaction where the human is assumed to yield to the robot. The same simulation capability could also be used to implement even more sophisticated selection strategies, for example, a cognitive equivalent to the minimax (or min-max) algorithm (von Neumann and Morgenstern 1944) with total distance traveled by the team as the cost function. Because our simple simulation was sufficient to meet the goal of this project, we did not pursue implementing any of these more sophisticated strategies.

TeamBot Performance

We ran our ACT-R/E system both on the physical robot in our lab and in a desktop simulator. (A video will be available.) In all runs, prior to the alarm, the TeamBot kept track of the whereabouts of its teammate by periodically looking at the teammate. This was indicated by a movement of the face shown on the robot's monitor to look at the teammate. When the alarm sounded, the TeamBot

would respond using either the self-centered or collaborative strategy.

Case 1 (A lack of teamwork): The default case was that both agents operated independently, i.e., the self-centered strategy, and each would go to their closest guard station. In this case, their decision making often resulted in both going to the same guard station, depending on the spatial situation at the time of the alarm.

Figure 3 is a trace of the locations of the human and robot during a run with each operating independently. The human began the run in the top of the diagram and the letters indicate the sequence of locations moving to the left: a, b, c, etc. The robot started at the bottom moving to the right. Both patrolled by moving counter-clockwise near the outer edges of the warehouse. At step "p" the alarm occurred and both started moving toward guard station #1 because it was closest to each of them at the time of the alarm. At time "u," the robot detected the conflict and started toward station #2. The conflict delayed completion of the alarm response task so that it ended at step "z."




Figure 3. Track of the human (starting at top) and TeamBot (starting at the bottom) 'manning' two guard stations (numbered squares) with a conflict. Both began at their respective "a" locations and patrolled counter-clockwise until the alarm occurred at "p." Then both moved toward guard station #1 until, at step "u," the TeamBot detects the conflict and goes to station #2.

Case 2 (The robot yields): A trace of the second case is shown in Figure 4. When the alarm occurred, the TeamBot evaluated what the human would do by simulating the human's decision making. It placed itself in the human's position and then used the same productions it would use to decide where to go, i.e., to the closest. Then knowing

where the human would go, it went to the other guard station thereby avoiding the conflict of both going to the same station. In this run, the team completed its task at step “x.”

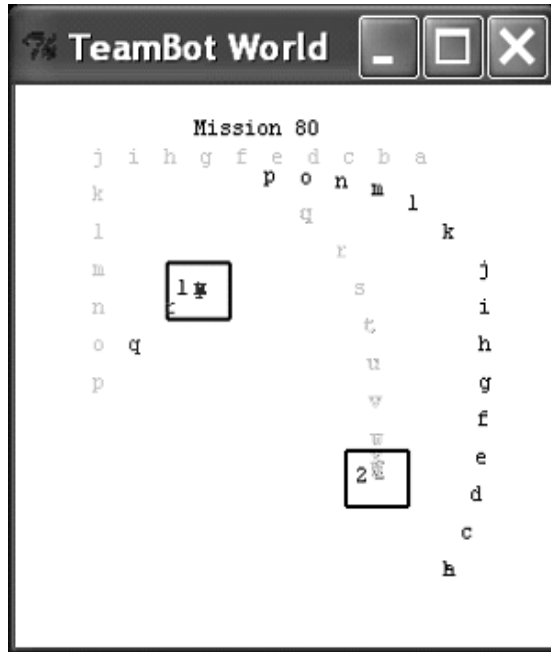


Figure 4. Track of the human (starting at top) and TeamBot (starting at the bottom) avoiding a conflict by simulating the human's decision making. Both began at their respective (a) locations and patrolled counter-clockwise until the alarm occurred at (p). The TeamBot then, in accordance with the shared plan, determined where the human would go by simulating the human's decision-making using its own knowledge and therefore immediately started toward the other station.

A series of desktop simulations of the system were run varying the starting positions of the TeamBot and its teammate along the top and bottom of the patrol area and with the human hardcoded to go to its closest station. They demonstrated that the performance with the TeamBot simulating the decision-making of its human teammate was significantly faster in achieving the goal of having ‘manning’ both guard stations after the alarm. Specifically, with 25 simulated runs each, when the TeamBot simulated the decision making of its teammate took 3.28 fewer steps than the system that did not, $t(27.7) = 8.1492$, $p < .001$ with the Welch correction for unequal variances. Experimental results are reported using desktop simulations; the actual robot's behavior was very similar.

We ran the same system on the iRobot B21r in our laboratory and demonstrated the successful embodiment of the system. A video of each case is available on our public website (www.nrl.navy.mil/aic/iss/) and our code is available from the authors.

Discussion

We claim to have achieved the principal goal in this project which was to show how a novel, scientifically-principled integration of simulation of a teammate within an embodied computational cognitive model can facilitate the teamwork aspects of human-robot interaction. Our robot reasoned about its teammate based on its own understanding of the situation and projected the decision-making of the human based on its own processes for appropriate teamwork behavior. The performance of the TeamBot demonstrates that integrating mental simulation can result in a more effective robotic teammate. We were able to incorporate mental simulation into a computational cognitive architecture (ACT-R) which itself was integrated with AI and robotics into an imbedded architecture, ACT-R/E. The result was a robotic teammate that supported the human-robot team performing statistically better in our warehouse security team scenario.

This result also supports *the representation hypothesis*. When the TeamBot used cognitive processes similar to a person's, i.e., mental simulation, it performed better as a teammate than when it did not and a conflict occurred and even with conflicts, it showed reasonable behavior.

Through the development of ACT-R/E, we have driven cognitive plausibility deeper into the system. In our previous work, an unmodified ACT-R computational cognitive modeling system worked through a middle layer to interact with the robotic systems (Kennedy et al. 2007). That system provided inputs to the cognitive model by inserting sensor information into the declarative memory and the system “remembered” the information. The new system, ACT-R/E, interfaces with the external world through buffers and modules in a more scientifically principled way and more consistently with the rest of the design of ACT-R, therefore making it a highly coherent, integrated system and more cognitively plausible. The achieved behavior is justifiably intelligent because of the underlying ACT-R architecture and its long history matching of human performance over a wide range of cognitive tasks.

The differences between the two tested cases in the warehouse scenario are, of course, sensitive to the distances between the objects in the scenario. We did not exhaustively test alternate configurations. In a more advanced scenario, communication between the team members or longer range detection of conflicts would reduce the difference between the two cases examined. We remain confident that there is a benefit of simulating the teammate's decision making avoids conflicts in any reasonable scenario setup.

Finally, this work demonstrated the value of implementing integrated architectures on mobile robots which can clearly test the effectiveness of the advancements.

Acknowledgements

Our work benefited from discussions with Dan Bothell and the vision work of Ben Fransen. This work was performed while the first author held a National Research Council Research Associateship Award and was partially supported by the Office of Naval Research under job order numbers N001406WZ20001 and N0001406WR20156. The views and conclusions contained in this document should not be interpreted as necessarily representing official policies, either expressed or implied, of the U.S. Navy.

References

- Anderson, J. R. 2007. *How Can the Human Mind Occur in the Physical Universe?* Oxford University Press.
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglas, S., Lebiere, C., and Qin, Y. 2004. An integrated theory of mind. *Psychological Review* 111(4): 1036-1060.
- Bello, P., and Cassimatis, N. L. 2006. Understanding other Minds: A Cognitive Approach. In *Proceedings of the Seventh International Conference on Cognitive Modeling*, 44-59. Trieste, IT: Edizioni Goliardiche.
- Cannon-Bowers, J. A., Salas, E., and Converse, S. 1993. Shared mental models in expert team decision making. In *Individual and group decision making*, 221-246. N. J. Castellan ed. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Cassimatis, N. L. 2005. Integrating cognitive models based on different computational methods. In *Proceedings of the XVII Annual Conference of the Cognitive Science Society*, 402-407. Mahwah, NJ: Lawrence Erlbaum Associates.
- Cooke, N. J., Salas, E., Cannon-Bowers, J. A., and Stout, R. 2000. Measuring team knowledge. *Human Factors* 42: 151-173.
- Falak-Ytter, T., Gredeback, G., and Von Hofsten, C. 2006. Infants predict other people's action goals. *Nature Neuroscience* 9: 878-879.
- Fransen, B. R., Morariu, V., Martinson, E., Blisard, S., Marge, M., Thomas, S., et al. 2007. *Using Vision, Acoustics, and Natural Language for Disambiguation*. Proceedings of the 2007 ACM Conference on Human-Robot Interaction, 73-80. Arlington, VA: ACM Press.
- Grosz, B. J., and Kraus, S. 1996. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2).
- Kaminka, G., Fidanboyly, M., Chang, A., and Velso, M. 2003. Learning the sequential coordinated behavior of teams from observation. In *RoboCup 2002: Robot Soccer World Cup IV*, LNAI 2752, 111-125. G. Kaminka, P. Lima and R. Royas eds. Fukuoka, Japan.
- Kaminka, G. A., and Tambe, M. 2000. Robust Agent Teams via Socially-Attentive Monitoring. *Journal of Artificial Intelligence Research* 12: 105-147.
- Kennedy, W. G., Bugajska, M. D., Marge, M., Fransen, B. R., Adams, W., Perzanowski, D., et al. 2007. Spatial Representation and Reasoning for Human-Robot Collaboration. In *Proceedings of the National Conference of the Association for the Advancement of Artificial Intelligence*, 1554-1559. Vancouver, BC: AAAI Press.
- Laird, J. E. 2001. It knows what you're going to do: adding anticipation to a quakebot. In *Proceedings of the Fifth International Conference on Autonomous Agents*. ACM.
- Mathieu, J. E., Heffner, T. S., Goodwin, G. F., Salas, E., and Cannon-Bowers, J. A. 2000. The influence of shared mental models on team process and performance. *Journal of Applied Psychology* 8(2): 273-283.
- Parke, F., and Waters, K. 1996. *Computer facial animation*. Natick, MA: AK Peters, Ltd.
- Rich, C., and Sidner, C. 1997. COLLAGEN: When agents collaborate with people. In *Proceedings of the First International Conference on Autonomous Agents*.
- Rosenbloom, P. S., Laird, J., and Newell, A. 1993. *The SOAR Papers*. Cambridge, MA: MIT Press.
- Schultz, A. C., Adams, W., and Yamauchi, B. 1999. Integrating exploration, localization, navigation, and planning with a common representation. *Autonomous Robots* 6: 293-308.
- Simmons, R., Adams, W., Atrash, A., Bugajska, M., Coblenz, M., MacMahon, M., et al. 2003. GRACE: an autonomous robot for the AAAI Robot challenge. *AI Magazine* 24: 51-72.
- Stone, P., and Velso, M. 2000. Multi-Agent systems: A Survey from a Machine Learning Perspective. *Autonomous Robots* 8(3): 345-383.
- Stout, R., Cannon-Bowers, J. A., Salas, E., and Milanovich, D. M. 1999. Planning, shared mental models, and coordinated performance: An empirical link is established. *Human Factors* 41: 61-71.
- Sycara, K. 1998. Multiagent Systems. *AI Magazine* 19: 79-92.
- Sycara, K., Paolucci, M., Giampapa, J., and van Velsen, M. 2003. The RETSINA multiagent infrastructure. *Autonomous Agents and Multi-agent Systems* 7(1).
- Sycara, K., and Sukthankar, G. 2006. *Literature Review of Teamwork Models*: Carnegie Mellon University.
- Tambe, M. 1997. Towards flexible teamwork. *Journal of Artificial Intelligence Research* 7: 83-124.
- Trafton, J. G., Bugajska, M. D., Fransen, B. R., and Ratwani, R. M. 2008. Integrating vision and audition within a cognitive architecture to track conversation. In *Proceedings of the 2008 ACM Conference on Human-Robot Interaction*, 201-208. Amsterdam, NL: ACM Press.
- Trafton, J. G., Cassimatis, N. L., Bugajska, M. D., Brock, D. P., Mintz, F. E., and Schultz, A. C. 2005. Enabling effective human-robot interaction using perspective-taking in robots. *IEEE Transactions on Systems, Man, and Cybernetics* 35(4): 460-470.
- von Neumann, J., and Morgenstern, O. 1944. *Theory of Games and Economic Behavior*. Princeton University Press.